

Modeling Spatial Extremes Using Normal Mean-Variance Mixtures

Zhongwei Zhang

CEMSE Division, King Abdullah University of Science and Technology
and

Raphaël Huser

CEMSE Division, King Abdullah University of Science and Technology
and

Thomas Opitz

Biostatistics and Spatial Processes, INRAE

and

Jennifer Wadsworth

Department of Mathematics and Statistics, Lancaster University

Abstract

Classical models for multivariate or spatial extremes are mainly based upon the asymptotically justified max-stable or generalized Pareto processes. These models are suitable when asymptotic dependence is present, i.e., the joint tail decays at the same rate as the marginal tail. However, recent environmental data applications suggest that asymptotic independence is equally important and, unfortunately, existing spatial models in this setting that are both flexible and can be fitted efficiently are scarce. Here, we propose a new spatial copula model based on the generalized hyperbolic distribution, which is a specific normal mean-variance mixture and is very popular in financial modeling. The tail properties of this distribution have been studied in the literature, but with contradictory results. It turns out that the proofs from the literature contain mistakes. We here give a corrected theoretical description of its tail dependence structure and then exploit the model to analyze a simulated dataset from the inverted Brown–Resnick process, hindcast significant wave height data in the North Sea, and wind gust data in the state of Oklahoma, USA. We demonstrate that our proposed model is flexible enough to capture the dependence structure not only in the tail but also in the bulk.

Keywords: Asymptotic independence; Extremal dependence modeling; Generalized hyperbolic distribution; Normal mean-variance mixtures; Spatial extremes.

1 Introduction

The statistical modeling of spatial extremes has gained significant interest in recent decades due to the increasing occurrence and sizes of natural extreme events, such as heat waves, heavy rainfall, and wildfires. When modeling spatial extremes, accurate inference for the marginal distribution at each site and a precise assessment of the dependence structure of extreme events among different sites are both needed; see Davison et al. (2012); Davison & Huser (2015) and Huser & Wadsworth (2020) for an overview. In this article we focus on modeling the spatial dependence of extreme events.

Classical models for spatial extremes are mainly based upon the asymptotically justified max-stable processes (de Haan, 1984; de Haan & Pereira, 2006) or generalized Pareto processes (Rootzén & Tajvidi, 2006; Ferreira & de Haan, 2014). The limiting dependence structures that arise in these models must be either asymptotically dependent (defined in Section 2.2), i.e., the joint tail decays at the same rate as the marginal tail, or for maxima, exactly independent. However, asymptotic independence appears to be equally important, as suggested by recent environmental data applications (Wadsworth & Tawn, 2012; Le et al., 2018). Unfortunately, existing models in this setting that are both flexible and can be fitted efficiently are scarce.

Existing models from the literature that can capture asymptotic independence include the Gaussian copula model (Bortot et al., 2000), inverted max-stable models (Wadsworth & Tawn, 2012), the Huser–Wadsworth model (Huser & Wadsworth, 2019), and the conditional extremes model (Wadsworth & Tawn, 2019). The Gaussian copula is the simplest model and has a restrictive dependence structure; inverted max-stable models present the same difficulties for inference as max-stable models since often only likelihoods based on lower-dimensional densities are available (Padoan et al., 2010; Castruccio et al., 2016); the Huser–Wadsworth model can capture both asymptotic dependence and independence but its distribution and (potentially censored) density functions rely on unidimensional integrals which have to be computed numerically, and this is computationally prohibitive in high dimensions, especially when computation of the multivariate normal distribution function is required; the conditional extremes model allows the change of asymptotic dependence class with distance between sites but it lacks an unconditional interpretation; see Huser & Wadsworth (2020) for a more detailed discussion.

Here we propose a flexible model based on an extension of scale mixtures. A random vector $\mathbf{X} \in \mathbb{R}^d$ is called a scale mixture if it has the representation $\mathbf{X} = R\mathbf{W}$, where R is a non-negative univariate random variable and $\mathbf{W} \in \mathbb{R}^d$ is a random vector. Scale mixtures provide a flexible family of distributions that can capture both asymptotic dependence and independence depending on the specification of the tail

of R and \mathbf{W} ; see Huser et al. (2017), Huser & Wadsworth (2019), Engelke et al. (2019). In practice \mathbf{W} is often taken as a Gaussian random vector and in this case \mathbf{X} is termed a Gaussian scale mixture. Natural alternatives and extensions of Gaussian scale mixtures include the Gaussian location mixtures (Krupskii et al., 2018), and the skew- t model (Morris et al., 2017), which is a specific Gaussian location-scale mixture. Here we focus on a different form of Gaussian location-scale mixtures, for which the term *normal mean-variance mixtures* has been coined in the literature. Specifically, \mathbf{X} is called a normal mean-variance mixture if it can be represented as

$$\mathbf{X} = \boldsymbol{\mu} + \gamma R + \sqrt{R}\mathbf{W}, \quad R \perp\!\!\!\perp \mathbf{W}, \quad (1)$$

where $\boldsymbol{\mu} \in \mathbb{R}^d$ is a location parameter vector, $\gamma \in \mathbb{R}^d$ regulates the skewness, $R \sim F_R$ is a non-negative mixing random variable, and $\mathbf{W} \sim \mathcal{N}_d(\mathbf{0}, \Sigma)$ is a Gaussian random vector with correlation matrix Σ . This provides a very flexible family of distributions and by properly choosing the parameters and mixing distribution, many well-known multivariate distributions can be obtained. Another benefit of normal mean-variance mixtures is the simple implementation of their conditional simulation, using their property of closedness under conditioning (Jamalizadeh & Balakrishnan, 2019), which is in clear contrast to max-stable or inverted max-stable models.

A prominent example of normal mean-variance mixtures is the generalized hyperbolic (GH) distribution, when the mixing distribution F_R is taken as generalized inverse Gaussian. The skew- t model (Morris et al., 2017) can be represented as $\mathbf{X} = \boldsymbol{\mu} + \gamma_c \mathbf{1}_d \sqrt{R}|Z| + \sqrt{R}\mathbf{W}$, $\gamma_c \in \mathbb{R}$, where $\mathbf{1}_d$ is a d -dimensional vector of 1s, R has an inverse Gamma distribution, which is a limiting case of the generalized inverse Gaussian distribution, Z is a standard normal random variable, \mathbf{W} remains a Gaussian random vector and R, Z, \mathbf{W} are mutually independent. Although the representations of this skew- t model and the GH distribution look similar, a major difference between them is the extra random variable Z in the representation of the skew- t model, and their tail dependence structures are significantly different. That is, the skew- t model is asymptotically dependent, and the GH distribution has more free parameters and is asymptotically independent except in one limiting case; see Section 2.3 for more details.

The GH distribution is very popular in financial modeling (Barndorff-Nielsen, 1997; Prause, 1999) thanks to its flexible univariate distribution and infinite divisibility; more details of this distribution are given in Section 2.1. Due to the popularity of this distribution in finance, its tail properties have been studied in the literature (Nolde, 2014; von Hammerstein, 2016), but surprisingly with contradictory results. By examining their proofs in detail, we have found that both of them contain subtle mistakes.

Here, we point out the mistakes in their proofs, which lead to the contradiction, and give a corrected description of the tail dependence of the GH distribution.

Based on this result, we propose to use the GH copula for spatial extremes, which has been less investigated compared with its flexible univariate distribution, and to our knowledge not yet exploited in environmental applications. This model provides flexible dependence structures with many subclasses and limiting models, such as the normal inverse Gaussian, hyperbolic, Gaussian and student t copulas. We apply the model using a full likelihood approach (thus, avoiding the computationally prohibitive censoring mechanism) to a simulated dataset from the inverted Brown–Resnick process, the hindcast significant wave height data considered in Wadsworth & Tawn (2012) and Huser & Wadsworth (2019), and wind gust data in the state of Oklahoma, USA, and demonstrate that our proposed model is flexible enough to capture the dependence structure not only in the tail but also in the bulk.

This paper is structured as follows. In Section 2 we introduce the multivariate generalized hyperbolic distribution and measures of tail dependence, and we then review existing results on the tail dependence properties of the GH distribution and present a corrected description. Section 3 introduces the copula-based likelihood inference, followed by two simulation studies, and Section 4 consists of two data applications. Section 5 concludes with a discussion.

2 Modeling

2.1 The Generalized Hyperbolic (GH) Distribution

The GH distribution is a specific normal mean-variance mixture with representation (1) and mixing distribution F_R specified to be generalized inverse Gaussian, denoted as $\text{GIG}(\lambda, \kappa, \psi)$. The $\text{GIG}(\lambda, \kappa, \psi)$ probability density function is

$$f_{\text{GIG}}(x) = \left(\frac{\psi}{\kappa}\right)^{\lambda/2} \frac{x^{\lambda-1}}{2K_\lambda(\sqrt{\kappa\psi})} \exp\left\{-\frac{1}{2}\left(\frac{\kappa}{x} + \psi x\right)\right\}, \quad x > 0,$$

where K_λ is the modified Bessel function of the second kind with index λ , and, for $\mathbf{x} \in \mathbb{R}^d$, the probability density function of a d -dimensional GH distribution is

$$f_{\text{GH}}(\mathbf{x}) = a_d \cdot \frac{K_{\lambda-d/2} \left[\sqrt{\{\kappa + (\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})\}} (\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \right] e^{(\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1} \boldsymbol{\gamma}}}{\left[\sqrt{\{\kappa + (\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})\}} (\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \right]^{d/2-\lambda}},$$

where

$$a_d = \frac{\psi^{\lambda/2} (\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma})^{d/2-\lambda}}{(2\pi)^{d/2} |\Sigma|^{1/2} \kappa^{\lambda/2} K_\lambda(\sqrt{\kappa\psi})},$$

Table 1: Some special subclasses and limiting cases of the GH distribution

Distribution	Parameter domain			
	λ	κ	ψ	$\boldsymbol{\gamma}$
Hyperbolic	$(d+1)/2$	> 0	> 0	$\in \mathbb{R}^d$
Normal inverse Gaussian (NIG)	$-1/2$	> 0	> 0	$\in \mathbb{R}^d$
Student- t with df degrees of freedom	$-\text{df}/2 < 0$	$\text{df} > 0$	0	$\mathbf{0}$
Cauchy	$-1/2$	1	0	$\mathbf{0}$

$\boldsymbol{\mu} \in \mathbb{R}^d$ is a location parameter vector, $\boldsymbol{\gamma} \in \mathbb{R}^d$ regulates the skewness, $\Sigma \in \mathbb{R}^{d \times d}$ is a positive definite dispersion matrix, and $\lambda, \kappa, \psi \in \mathbb{R}$ control the shape of the mixing GIG distribution. We denote this GH distribution by $\mathbf{X} \sim \text{GH}_d(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \boldsymbol{\mu}, \Sigma)$. Both in the GIG and GH distributions, the admissible parameter values are $\lambda < 0, \kappa > 0, \psi \geq 0$ or $\lambda = 0, \kappa > 0, \psi > 0$ or $\lambda > 0, \kappa \geq 0, \psi > 0$. Note that $\psi = 0$ and $\kappa = 0$ should be understood as two limiting cases.

One important property of the GH distribution is that it is closed under marginalization, conditioning and linear transformations. Specifically, if $\mathbf{X} \sim \text{GH}_d(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \boldsymbol{\mu}, \Sigma)$ and $\mathbf{Y} = B\mathbf{X} + \mathbf{b}$, where $B \in \mathbb{R}^{k \times d}$ and $\mathbf{b} \in \mathbb{R}^k$, then $\mathbf{Y} \sim \text{GH}_k(\lambda, \kappa, \psi, B\boldsymbol{\gamma}, B\boldsymbol{\mu} + \mathbf{b}, B\Sigma B^\top)$. This distribution is a rich family with many special subclasses and limiting cases; see Table 1 for some examples and the corresponding admissible parameter domains. Note that, for the GH, hyperbolic, and NIG distributions the dispersion matrix Σ needs to be a correlation matrix for identifiability reasons. This distribution was first introduced in Barndorff-Nielsen (1977) to model sand sizes and has become very popular in modeling turbulence and returns of financial assets; see Barndorff-Nielsen (1997), Prause (1999), and Schmidt (2003). It is also worthwhile to mention that the GH distribution is elliptical if and only if $\boldsymbol{\gamma} = \mathbf{0}$; see Corollary 3 in von Hammerstein (2016). More details on its properties can be found in Blaesild & Jensen (1981), Prause (1999) and McNeil et al. (2005).

There are several equivalent parametrizations for the GH distribution. The most common one is to parametrize it as $\text{GH}(\lambda, \alpha, \boldsymbol{\mu}, \Delta, \delta, \boldsymbol{\beta})$, where we let $\Delta = |\Sigma|^{-1/d}\Sigma$, $\boldsymbol{\beta} = \Sigma^{-1}\boldsymbol{\gamma}$, $\delta = \sqrt{|\Sigma|^{1/d}\kappa}$, and $\alpha = \sqrt{|\Sigma|^{-1/d}(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1}\boldsymbol{\gamma})}$. This parametrization was used in Barndorff-Nielsen (1977), but it does not give the nice property that the important parameters α and δ are invariant under linear transformations. Therefore, we here instead adopt the parametrization $\text{GH}_d(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \boldsymbol{\mu}, \Sigma)$.

2.2 Measures of Tail Dependence

Coles et al. (1999) proposed the tail dependence coefficient χ , providing it exists, to measure the extremal dependence between random variables X_1 and X_2 . It may be defined through the limit

$$\chi_u := \frac{\Pr\{F_1(X_1) > u, F_2(X_2) > u\}}{\Pr\{F_1(X_1) > u\}} \rightarrow \chi \text{ as } u \rightarrow 1, \quad (2)$$

where F_1 and F_2 are the marginal distribution functions of X_1 and X_2 , respectively, and are assumed to be continuous without loss of generality. The lower tail dependence coefficient can be obtained via reflection. Throughout this work we focus on the upper tail dependence. The vector $(X_1, X_2)^\top$ is termed asymptotically independent if $\chi = 0$ and asymptotically dependent if $\chi > 0$.

When asymptotic independence is present, one useful quantity to measure the residual dependence at pre-asymptotic levels is the residual tail dependence coefficient η (Ledford & Tawn, 1996), which is defined by the following asymptotic expansion of the joint tail (providing it exists):

$$\Pr\{F_1(X_1) > u, F_2(X_2) > u\} \sim \mathcal{L}\{(1-u)^{-1}\}(1-u)^{1/\eta} \text{ as } u \rightarrow 1, \quad (3)$$

where \mathcal{L} is a slowly varying function, i.e., $\mathcal{L}(tx)/\mathcal{L}(x) \rightarrow 1$ as $x \rightarrow \infty$ for all $t > 0$, and $f(x) \sim g(x)$ as $x \rightarrow x_0$ means that, for g non-zero in a neighbourhood of x_0 , $f(x)/g(x) \rightarrow 1$ as $x \rightarrow x_0$. The coefficient $\eta \in (0, 1]$ determines the joint tail decay rate. A reformulation of (3) implies that η can be alternatively defined through the limit

$$\eta_u := \frac{\log \Pr\{F_1(X_1) > u\}}{\log \Pr\{F_1(X_1) > u, F_2(X_2) > u\}} \rightarrow \eta \text{ as } u \rightarrow 1, \quad (4)$$

and this definition allows one to estimate η with η_u for u close to 1.

The relationship between these two above coefficients is $\eta < 1 \implies \chi = 0$ and $\chi > 0 \implies \eta = 1$. The converse of this relationship does not necessarily hold; see Proposition 4 in Manner & Segers (2011) and an example with $\eta = 1, \chi = 0$ in Huser & Wadsworth (2019). We say that $(X_1, X_2)^\top$ is (a) positively associated if $1/2 < \eta \leq 1$; (b) near-independent if $\eta = 1/2$; (c) negatively associated if $0 < \eta < 1/2$. The bivariate measures χ and η can be easily extended to dimension $d > 2$ by replacing $\Pr\{F_1(X_1) > u, F_2(X_2) > u\}$ in (2) and (4) with $\Pr\{F_1(X_1) > u, \dots, F_d(X_d) > u\}$. As a result, the pair of coefficients (χ, η) is able to measure the tail dependence strength across asymptotic dependence and independence classes, and it is often used for assessing model fit.

2.3 Tail Dependence of the GH Distribution

Recall that the GH distribution is elliptical if and only if $\boldsymbol{\gamma} = \mathbf{0}$. Schlueter & Fischer (2012) have proved that the elliptical GH distribution with dispersion matrix $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, 0 \leq \rho < 1$, excluding the limiting case $\psi = 0$, is asymptotically independent and has residual tail dependence coefficient $\eta = \sqrt{(1 + \rho)/2}$; see also Huser et al. (2017). Nolde (2014) considered a more general class and claimed that the whole GH distribution family with dispersion matrix $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, |\rho| \neq 1$, has residual tail dependence coefficient $\sqrt{(1 + \rho)/2}$, which implies that the skewness parameter $\boldsymbol{\gamma}$ has no influence on the residual tail dependence coefficient. However, von Hammerstein (2016) used a different approach and claimed that the GH distribution can be asymptotically dependent or independent depending on the choice of the parameters, although without giving the residual tail dependence coefficient when asymptotic independence is present.

After carefully investigating where this contradiction comes from, we found that both of their proofs have mistakes which lead to incorrect conclusions. We detail their mistakes and a corrected version of von Hammerstein's proof in the Supplementary Material. We now give a corrected description of the tail dependence property of the GH distribution, using a similar approach to Nolde (2014). In contrast with the often complicated calculation of the residual tail dependence coefficient by its definition, Nolde (2014) provided a geometric interpretation of this coefficient, which leads to simple and intuitive computations of this coefficient for a variety of distributions. This is particularly the case when joint densities are easier to compute than joint distribution or survival functions. We first recall some definitions and theorems from Nolde (2014) which will be used in our proof.

Consider a sequence of independent and identically distributed random vectors $\mathbf{Z}_1, \mathbf{Z}_2, \dots$ on \mathbb{R}^2 . Let $N_n := \{\mathbf{Z}_1/r_n, \dots, \mathbf{Z}_n/r_n\}$ denote an n -point sample cloud with scaling constants $r_n > 0, r_n \rightarrow \infty$ as $n \rightarrow \infty$ and let $N_n(A)$ be the number of points of N_n contained in the Borel set $A \in \mathbb{R}^d$, i.e. $N_n(A) = \sum_{i=1}^n 1_A(\mathbf{Z}_i/r_n)$. Let D be a compact set in \mathbb{R}^2 . Then D is called a *limit set* of the sample cloud N_n as $n \rightarrow \infty$ if (i) $\Pr\{N_n(U^c) > 0\} \rightarrow 0$ for open sets U containing D , and U^c denotes the complement of U , and (ii) $\Pr\{N_n(\mathbf{p} + \varepsilon B) > m\} \rightarrow 1$ for all $m \geq 1, \varepsilon > 0, \mathbf{p} \in D$, where B denotes the Euclidean unit ball. If every ray from the origin intersects the boundary of a given set D in a single point, then D can be characterized by a continuous *gauge function* n_D . That is, $n_D : \mathbb{R}^2 \rightarrow [0, \infty)$ is a homogeneous function of degree one, i.e. $n_D(t\mathbf{x}) = tn_D(\mathbf{x})$ for all $t > 0$, and $D = \{\mathbf{x} \in \mathbb{R}^2 : n_D(\mathbf{x}) < 1\}$. A set D is called *star shaped* if $\mathbf{x} \in D$ implies $t\mathbf{x} \in D$ for all $t \in (0, 1)$.

A measurable function $\Psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is *regularly varying* at infinity with exponent $a \in \mathbb{R}$ if for

$x > 0$, $\lim_{t \rightarrow \infty} \Psi(tx)/\Psi(x) = x^a$. Roughly speaking, regularly varying functions behave asymptotically like power functions. When $a = 0$, the function Ψ is slowly varying.

Lemma 2.1 (Theorem 2.1 in Nolde (2014)). *Let $\{(X_i, Y_i)^\top, i \geq 1\}$ be independent random vectors in \mathbb{R}^2 from a distribution G with marginal survival functions $1 - G_i(s) = e^{-\Psi_i(s)}$, $s > 0$, where Ψ_1 and Ψ_2 are regularly varying at infinity with the same exponent $a > 0$. Suppose there exists a scaling sequence $r_n > 0$ such that the n -point sample cloud*

$$N_n = \{(X_1/r_n, Y_1/r_n)^\top, \dots, (X_n/r_n, Y_n/r_n)^\top\}$$

converges onto set D whose interior is bounded, open, and star-shaped as $n \rightarrow \infty$. Let $\mathbf{q} = (q_1, q_2)^\top = \sup D$ denote the coordinate-wise supremum of D . Then D is a star-shaped subset of $(-\infty, \mathbf{q}]$. Define

$$r_D = \min\{r \geq 0 : D \cap ((rq_1, \infty) \times (rq_2, \infty)) = \emptyset\}.$$

If $\mathbf{q} \notin D$, then $r_D < 1$, X_1 and Y_1 are asymptotically independent and the residual dependence coefficient is $\eta = r_D^a$.

The definition of r_D implies that, if D is convex, $r_D \mathbf{q}$ lies on the boundary of D and thus $n_D(r_D \mathbf{q}) = 1$. Using the homogeneity of n_D and Lemma 2.1 one can get $\eta = n_D(\mathbf{q})^{-a}$; see also Nolde & Wadsworth (2020) for further details. In practice, a random vector is often described by its multivariate probability density function. The relation between the density function and its limit set is given in the following lemma and thereby a convenient way to calculate the residual tail dependence coefficient η is given.

Lemma 2.2 (Proposition 3.1 in Nolde (2014)). *Let $\mathbf{Z}_1, \mathbf{Z}_2, \dots$ be independent and identically distributed random vectors with a continuous positive density g on \mathbb{R}^d . Let $\gamma := -\log g$. Suppose there exists a function ℓ on \mathbb{R}^d which is positive outside a bounded set, and a nonzero vector \mathbf{v} such that*

$$\frac{\gamma(r_n \mathbf{u}_n)}{\gamma(r_n \mathbf{v})} \rightarrow \ell(\mathbf{u}), \text{ for some } r_n \rightarrow \infty, \text{ and any } \mathbf{u}_n \rightarrow \mathbf{u} \in \mathbb{R}^d.$$

Then the sequence of sample clouds $N_n = \{\mathbf{Z}_1/r_n, \dots, \mathbf{Z}_n/r_n\}$ converges as $n \rightarrow \infty$ onto a limit set D whose gauge function is given by $n_D = \ell^\theta$ for some positive constant θ .

For details about the proof of Lemma 2.1 and Lemma 2.2 and further background knowledge, we refer to Nolde (2014) and references therein. We are now ready to use these two results to show that the GH distribution, except in the limiting case $\psi = 0$, is asymptotically independent and we give its residual dependence coefficient.

Proposition 2.3. Let $(X_1, X_2)^\top \sim GH_2(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \boldsymbol{\mu}, \Sigma)$ with $\boldsymbol{\gamma} = (\gamma_1, \gamma_2)^\top$, $\boldsymbol{\mu} = (\mu_1, \mu_2)^\top$, and $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, $|\rho| \neq 1$. If $\psi > 0$, then X_1, X_2 are asymptotically independent and the residual dependence coefficient is

$$\eta = \frac{1 - \rho^2}{\sqrt{\{\psi(1 - \rho^2) + \gamma_1^2 - 2\rho\gamma_1\gamma_2 + \gamma_2^2\}(m_1^2 - 2\rho m_1 m_2 + m_2^2) - m_1(\gamma_1 - \rho\gamma_2) - m_2(\gamma_2 - \rho\gamma_1)}},$$

where

$$m_1 = \frac{\gamma_1 + \sqrt{\psi + \gamma_1^2}}{\psi}, \quad m_2 = \frac{\gamma_2 + \sqrt{\psi + \gamma_2^2}}{\psi}.$$

Proof. Note that both the tail dependence coefficient and the residual dependence coefficient are copula properties, i.e., they are invariant under strictly increasing marginal transformations. Hence, for simplicity we assume $\boldsymbol{\mu} = \mathbf{0}$ and study the tail dependence of $(X_1, X_2)^\top \sim GH_2(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \mathbf{0}, \Sigma)$. Let f be the probability density function of $(X_1, X_2)^\top$. Using the asymptotic property of the Bessel function (see formula (9.7.2) in Abramowitz & Stegun (1972)), we know that $K_\lambda(x) = \sqrt{\frac{\pi}{2x}} \exp(-x) + o(1/\sqrt{x})$ as $x \rightarrow \infty$. Then, for $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$ and $s \rightarrow \infty$, we have

$$\begin{aligned} \frac{-\log\{f(s\mathbf{u})\}}{-\log\{f(s\mathbf{v})\}} &\sim \frac{\sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma})(\kappa + s^2 \mathbf{u}^\top \Sigma^{-1} \mathbf{u})} - s \mathbf{u}^\top \Sigma^{-1} \boldsymbol{\gamma}}{\sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma})(\kappa + s^2 \mathbf{v}^\top \Sigma^{-1} \mathbf{v})} - s \mathbf{v}^\top \Sigma^{-1} \boldsymbol{\gamma}} \\ &\sim \frac{\sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \mathbf{u}^\top \Sigma^{-1} \mathbf{u}} - \mathbf{u}^\top \Sigma^{-1} \boldsymbol{\gamma}}{\sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \mathbf{v}^\top \Sigma^{-1} \mathbf{v}} - \mathbf{v}^\top \Sigma^{-1} \boldsymbol{\gamma}}. \end{aligned}$$

For fixed $\mathbf{v} \neq \mathbf{0}$, let $h := \sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \mathbf{v}^\top \Sigma^{-1} \mathbf{v}} - \mathbf{v}^\top \Sigma^{-1} \boldsymbol{\gamma}$. Then, consider the set D :

$$D = \{\mathbf{u} \in \mathbb{R}^2 : (\sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \mathbf{u}^\top \Sigma^{-1} \mathbf{u}} - \mathbf{u}^\top \Sigma^{-1} \boldsymbol{\gamma})/h \leq 1\}.$$

Note that $(\sqrt{(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \mathbf{u}^\top \Sigma^{-1} \mathbf{u}} - \mathbf{u}^\top \Sigma^{-1} \boldsymbol{\gamma})/h \leq 1$ is equivalent to

$$(\psi + \boldsymbol{\gamma}^\top \Sigma^{-1} \boldsymbol{\gamma}) \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \leq (h + \mathbf{u}^\top \Sigma^{-1} \boldsymbol{\gamma})^2,$$

which can be simplified to

$$(\psi + \gamma_2^2)u_1^2 + (\psi + \gamma_1^2)u_2^2 - 2(\rho\psi + \gamma_1\gamma_2)u_1u_2 - 2h\{(\gamma_1 - \rho\gamma_2)u_1 + (\gamma_2 - \rho\gamma_1)u_2\} \leq h^2(1 - \rho^2). \quad (5)$$

Since $\psi > 0, \rho \neq 1$, we have

$$\begin{aligned} \Delta &= 4(\rho\psi + \gamma_1\gamma_2)^2 - 4(\psi + \gamma_2^2)(\psi + \gamma_1^2) \\ &= -4(1 - \rho^2)\psi^2 - 4\psi(\gamma_1^2 + \gamma_2^2 - 2\rho\gamma_1\gamma_2) \\ &< 0. \end{aligned}$$

This implies that D is an ellipse, and thus convex. By Lemma 2.2, D is the limit set and its associated gauge function is $n_D(\mathbf{u}) = \{\sqrt{(\psi + \boldsymbol{\gamma}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}) \mathbf{u}^\top \boldsymbol{\Sigma}^{-1} \mathbf{u} - \mathbf{u}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}}\}/h$. Furthermore, note that $n_D(\mathbf{0}) \leq 1$, which implies that $\mathbf{0}$ is contained in D . Hence, D is bounded and star shaped. Now it remains to find the coordinatewise supremum of D .

Note that the inequality (5) can be written as

$$(\psi + \gamma_2^2)u_1^2 - 2[(\rho\psi + \gamma_1\gamma_2)u_2 + h(\gamma_1 - \rho\gamma_2)]u_1 + (\psi + \gamma_1^2)u_2^2 - 2h(\gamma_2 - \rho\gamma_1)u_2 \leq h^2(1 - \rho^2).$$

Denote by $\mathbf{q} = (u_1^*, u_2^*)^\top$ the coordinatewise supremum of D , then we have

$$(\psi + \gamma_1^2)u_2^{*2} - 2h(\gamma_2 - \rho\gamma_1)u_2^* = h^2(1 - \rho^2) + \frac{\{(\rho\psi + \gamma_1\gamma_2)u_2^* + h(\gamma_1 - \rho\gamma_2)\}^2}{\psi + \gamma_2^2},$$

which can be simplified to

$$\psi u_2^{*2} - 2h\gamma_2 u_2^* - h^2 = 0.$$

Hence,

$$u_2^* = \max\left(\frac{2h\gamma_2 + \sqrt{4h^2\gamma_2^2 + 4h^2\psi}}{2\psi}, \frac{2h\gamma_2 - \sqrt{4h^2\gamma_2^2 + 4h^2\psi}}{2\psi}\right) = \frac{h(\gamma_2 + \sqrt{h^2\gamma_2^2 + h^2\psi})}{\psi}.$$

Similarly, we get $u_1^* = h(\gamma_1 + \sqrt{h^2\gamma_1^2 + h^2\psi})/\psi$. Since $(X_1, X_2)^\top \sim \text{GH}_2(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \mathbf{0}, \boldsymbol{\Sigma})$, we know that $X_i \sim \text{GH}_2(\lambda, \kappa, \psi, \gamma_i, 0, 1)$, $i = 1, 2$. Denote the marginal density function and distribution function of X_i , $i = 1, 2$ as f_{X_i} and F_{X_i} , respectively, and write $1 - F_{X_i}(x) = e^{-\Psi_i(x)}$. Using the asymptotic relation $K_\lambda(x) \sim \sqrt{\pi/(2x)}e^{-x}$, $x \rightarrow \infty$ we get

$$f_{X_i}(x) \sim c_i x^{\lambda-1} \exp\left\{-\left(\sqrt{\psi + \gamma_i^2} - \gamma_i\right)x\right\} \text{ as } x \rightarrow \infty,$$

where c_i is a constant. As $\psi > 0$ we know that $\sqrt{\psi + \gamma_i^2} - \gamma_i > 0$. By Proposition 2 in von Hammerstein (2016), we have

$$1 - F_{X_i}(x) \sim \frac{c_i x^{\lambda-1}}{\sqrt{\psi + \gamma_i^2} - \gamma_i} \exp\left\{-\left(\sqrt{\psi + \gamma_i^2} - \gamma_i\right)x\right\} \text{ as } x \rightarrow \infty.$$

Hence, the functions Ψ_i , $i = 1, 2$ are both regularly varying at infinity with the same exponent $a = 1$.

As the supremum point \mathbf{q} of an ellipse D satisfies $\mathbf{q} \notin D$, by Lemma 2.1 and Lemma 2.2, X_1 and X_2 are asymptotically independent and the residual dependence coefficient is

$$\begin{aligned} \eta &= n_D(\mathbf{q})^{-a} \\ &= [\{\sqrt{(\psi + \boldsymbol{\gamma}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}) \mathbf{q}^\top \boldsymbol{\Sigma}^{-1} \mathbf{q} - \mathbf{q}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}}\}/h]^{-1} \\ &= \frac{1 - \rho^2}{\sqrt{\{\psi(1 - \rho^2) + \gamma_1^2 - 2\rho\gamma_1\gamma_2 + \gamma_2^2\}(m_1^2 - 2\rho m_1 m_2 + m_2^2) - m_1(\gamma_1 - \rho\gamma_2) - m_2(\gamma_2 - \rho\gamma_1)}}, \end{aligned}$$

where

$$m_1 = \frac{\gamma_1 + \sqrt{\psi + \gamma_1^2}}{\psi}, \quad m_2 = \frac{\gamma_2 + \sqrt{\psi + \gamma_2^2}}{\psi}.$$

□

When $\rho = 0$, which corresponds to independence between the components of the Gaussian random vector \mathbf{W} and implies that the dependence between X_1 and X_2 is fully specified by the mixing variable R , we have $\eta = \{(\psi + \gamma_1^2 + \gamma_2^2)^{1/2}(m_1^2 + m_2^2)^{1/2} - m_1\gamma_1 - m_2\gamma_2\}^{-1}$, where m_1, m_2 are the same as above. One further observation is that when $\psi \rightarrow \infty$, which means the mixing variable R has a very light tail, we have $\eta = \sqrt{(1 + \rho)/2}$. In this case, we obtain the same η as the elliptical GH distribution, i.e., when $\boldsymbol{\gamma} = \mathbf{0}$, and it equals the square root of the residual dependence coefficient of a bivariate Gaussian random vector with correlation ρ .

It is also important to note that when $\psi = 0, \boldsymbol{\gamma} \neq \mathbf{0}$, the sample cloud does not converge onto a bounded set and thus the method we used above fails. When $\psi = 0, \boldsymbol{\gamma} = \mathbf{0}$, the GH distribution reduces to the (asymptotically dependent) Student- t distribution and the tail dependence coefficient is known from the literature; see Embrechts et al. (2001).

Figure 1 illustrates the flexibility in extremal dependence structures of the GH distribution, by plotting the bivariate χ_u, η_u and their limits as a function of distance between the pairs, for a range of values of $u \in [0.9, 1)$, and $\lambda = -0.5, \kappa = \psi = 1, \boldsymbol{\gamma} = \mathbf{0}$, Σ being a correlation matrix determined by a powered exponential correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2) = \exp\{-\|\mathbf{s}_1 - \mathbf{s}_2\|/0.6\}^{3/2}$, where $\|\mathbf{s}_1 - \mathbf{s}_2\|$ is the Euclidean distance between sites \mathbf{s}_1 and \mathbf{s}_2 . The results indicate the slow convergence of χ_u and η_u to their limits, and at any observable levels, including $u = 0.99$, nonnegligible dependence may exist (χ_u could be as large as 0.4 or 0.5 for pairs at short distances), but the pairs are asymptotically independent. This emphasizes again the need for asymptotically independent models or models that can capture both asymptotic independence and dependence.

3 Copula-Based Inference

3.1 Copula Model and Likelihood Inference

As here we focus on modeling the spatial dependence, we take a copula approach. A copula is a multivariate distribution function with standard uniform margins. Thanks to Sklar's theorem (Sklar, 1959), for any continuous multi-dimensional distribution function there is a unique copula associated




Figure 1: Bivariate χ_u , η_u for $u = 0.9, 0.95, 0.99$ and their limits as $u \rightarrow 1$ of the GH distribution for $\lambda = -0.5, \kappa = \psi = 1, \gamma = \mathbf{0}$, and Σ being a correlation matrix determined by a powered exponential correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2) = \exp\{-\|\mathbf{s}_1 - \mathbf{s}_2\|/0.6\}^{3/2}$, where $\|\mathbf{s}_1 - \mathbf{s}_2\|$ is the Euclidean distance between sites \mathbf{s}_1 and \mathbf{s}_2 . For $u < 1$, χ_u and η_u are calculated numerically using their definition (2) and (4), and their limits are calculated using the results in Proposition 2.3.

with it. This is the case for the multivariate GH distribution and we call the associated copula the GH copula. Specifically, suppose $(X_1, \dots, X_d)^\top$ follow a d -dimensional GH distribution F , which implies that their marginal distribution functions F_1, \dots, F_d are univariate GH distributions, then the GH copula is defined by

$$C(\mathbf{u}) = \Pr\{F_1(X_1) \leq u_1, \dots, F_d(X_d) \leq u_d\} = F\{F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)\}, \quad \mathbf{u} \in [0, 1]^d,$$

where $F_i^{-1}, i = 1, \dots, d$, is the inverse of F_i . Its density function thus can be derived easily as

$$c(\mathbf{u}) = \frac{\partial^d}{\partial u_1 \dots \partial u_d} C(\mathbf{u}) = \frac{f\{F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)\}}{f_1\{F_1^{-1}(u_1)\} \dots f_d\{F_d^{-1}(u_d)\}},$$

where f is the joint density function and f_1, \dots, f_d are the marginal GH density functions of X_1, \dots, X_d , respectively.

Due to the closed-form density of the multivariate GH distribution, inference based on the full likelihood is feasible. Here both in the simulation study and data applications, we fit the GH copula model to the whole dataset, rather than using a censored likelihood approach as Wadsworth & Tawn

(2012). By doing so, we gain two major benefits: (i) we can avoid the computational issues associated with the censored likelihood, which requires calculations of potentially high-dimensional distribution functions and that often have to be evaluated numerically; (ii) we can model dependence both in the bulk and in the tail. However, we also point out that an obvious disadvantage of taking such a full likelihood approach is that the extremal dependence structure would not be as well captured as using a censored likelihood approach. Here we choose the full likelihood approach because the GH copula model is highly flexible, more so than most other spatial models. This means that it could be a reasonable approach, and the simulation study will support this.

We leverage the R (R Core Team, 2020) package *Rcpp* to evaluate the GH density function with low-level language C++ and the package *ghyp* to compute the quantiles of the GH distribution, which uses splines to interpolate the univariate GH distribution function and then finds the root efficiently with the `uniroot` function in R. These “tricks” have improved the efficiency of the algorithm and, as an illustration, the computational time for fitting the GH copula to a dataset with 100 sites and 2000 observations at each site is less than three hours on a laptop with 8 GB memory and 2.3 GHz Intel Core i5 processor.

3.2 Simulation Study

Simulation from the multivariate GH distribution is straightforward due to its stochastic representation (1). Specifically, independently sampling $R \sim \text{GIG}(\lambda, \kappa, \psi)$ and $\mathbf{W} \sim \mathcal{N}_d(\mathbf{0}, \Sigma)$ and plugging them into formula (1) yields exact samples from the distribution $\text{GH}_d(\lambda, \kappa, \psi, \boldsymbol{\gamma}, \boldsymbol{\mu}, \Sigma)$. Sampling from a multivariate normal distribution is easy and sampling from a generalized inverse Gaussian distribution is feasible using the algorithm in Dagpunar (1989), which is implemented in the R package *ghyp*. Once a sample is generated from the multivariate GH distribution, one can obtain a sample for the GH copula by transforming the marginals to uniform distribution, using the probability integral transform.

To reduce the number of parameters in the GH copula model, in both the simulation study and data application, we here assume that $\boldsymbol{\gamma} = \gamma_c \mathbf{1}_d$, $\gamma_c \in \mathbb{R}$ and Σ is a correlation matrix determined by a powered exponential correlation function $\rho(\mathbf{s}_1, \mathbf{s}_2) = \exp\{-\|\mathbf{s}_1 - \mathbf{s}_2\|/\zeta\}^\nu$, where $\|\mathbf{s}_1 - \mathbf{s}_2\|$ is the Euclidean distance between sites \mathbf{s}_1 and \mathbf{s}_2 , $\zeta > 0$ is the range parameter, and $\nu \in (0, 2]$ controls the smoothness of the realized random fields. In the first simulation study we consider 25 sites on the uniform grid $\{0, 0.25, 0.5, 0.75, 1\}^2$, and set $\lambda = \kappa = \psi = \gamma_c = 1$, $\zeta = 0.2$ or 0.7 (short or long range dependence), $\nu = 0.5$ or 1.5 (rough or smooth random field). Then we generate $n = 250, 500$ or 1000

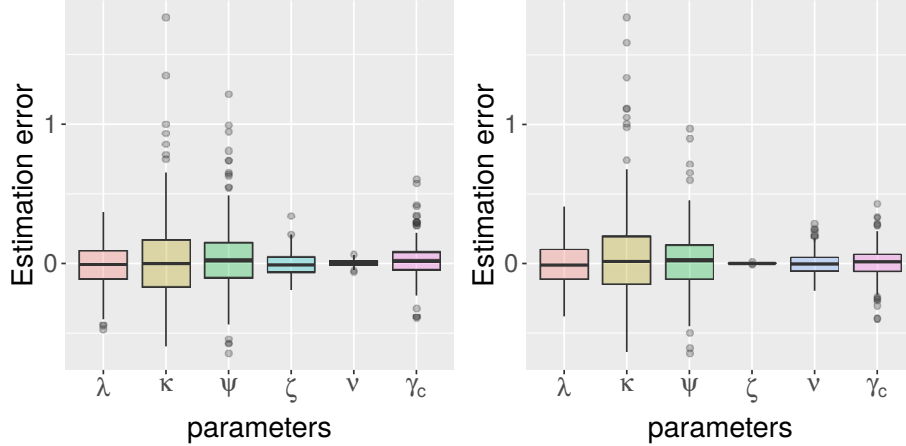


Figure 2: Boxplots of estimation errors of all the parameters for different scenarios in the first simulation study. Left panel: scenario $\zeta = 0.7, \nu = 0.5, n = 500$; right panel: scenario $\zeta = 0.2, \nu = 1.5, n = 500$.

replicates from the GH copula at these sites and fit the GH copula to the simulated dataset of each scenario.

We run the experiment for 300 times and Figure 2 depicts the boxplots of estimation errors (i.e., differences between parameter estimates and the true parameters) for all of the six parameters when $\zeta = 0.7, \nu = 0.5, n = 500$ and $\zeta = 0.2, \nu = 1.5, n = 500$. The results indicate that the spatial range parameter ζ and smoothness parameter ν are relatively easier to estimate than the parameters λ, κ, ψ , which determine the shape of the mixing GIG distribution, and the skewness parameter γ_c . We also observe that changing the values of ζ and ν does not affect the estimation of other parameters, but the larger value of ζ , which corresponds to stronger spatial dependence between sites, leads to smaller variability of the estimate of ν , and larger value of ν also yields smaller variability of the estimate of ζ . One further observation is that when we increase the number of replicates n , the variability of the estimates of all parameters decreases, roughly at a rate of \sqrt{n} as expected, as in shown in Figure 3.

In the second simulation study we consider 100 sites on the grid $\{0, 1/9, 2/9, \dots, 8/9, 1\}^2$ and generate 300 replicates from the inverted Brown–Resnick process (Kablichko et al., 2009; Wadsworth & Tawn, 2012) using the extremal functions approach of Dombry et al. (2016), with the variogram of the Brown–Resnick process defined as $\gamma(\mathbf{h}) = 2\|\mathbf{h}\|/\zeta$, where $\zeta = 0.3$. This model is known to be asymptotically independent with bivariate residual tail dependence coefficient function $\eta(\mathbf{h}) = 1/[2\Phi\{\sqrt{\gamma(\mathbf{h})}/2\}]$. To illustrate the flexibility of the GH distribution in a misspecified setting, we then fit the GH copula and the normal inverse Gaussian (NIG) copula, which is a subclass of the GH copula when $\lambda = -1/2$, to the




Figure 3: Boxplots of estimation errors of all the parameters for the scenario $\zeta = 0.7, \nu = 1.5$ with different number of replicates n in the first simulation study.

whole simulated dataset. Figure 4 displays the true residual tail dependence coefficient of the inverted Brown–Resnick process, its counterpart of the fitted GH copula and NIG copula model, and empirical estimates of η_u , as defined in (4), based on the simulated dataset. The results show that, despite fitting the GH copula to the entire dataset without censoring low non-extreme values, the GH copula captures the residual tail dependence quite well for the simulated dataset, and the shape of the fitted residual tail dependence coefficient function, with respect to distance, resembles its true function very closely, although a slight bias exists as expected.

4 Environmental Applications

4.1 Wave Height Example

We first consider the hindcast dataset of significant wave heights analyzed by Wadsworth & Tawn (2012) and Huser & Wadsworth (2019). This dataset contains eight observations per day over a period of 31 years, at 50 spatial locations in the North Sea. To reduce temporal dependence and ease the computational burden, only one observation per day in the winter months (December, January, and February) at




Figure 4: Estimated bivariate residual tail dependence coefficients in the second simulation study. Black line: true η of the inverted Brown–Resnick process; red line: theoretical η of the fitted normal inverse Gaussian copula, calculated using Proposition 2.3; blue line: theoretical η of the fitted GH copula, calculated using Proposition 2.3; green diamond: empirical tail dependence coefficient η_u with $u = 0.95$.

a subset of 20 sites has been used to fit the Huser–Wadsworth model from Huser & Wadsworth (2019). Here, we consider the same setting, although the GH copula or its subclasses are computationally more efficient to fit due to the closed-form density of the GH distribution.

The margins at each site are transformed to standard uniform following the semiparametric approach of Coles & Tawn (1991). Then, three different models are fitted to the standardized dataset, namely the Huser–Wadsworth model, the Gaussian copula model and the GH copula model. To illustrate the flexibility of our proposed GH model to capture the upper tail dependence in a disadvantageous situation, we first use a censored likelihood scheme to fit the Huser–Wadsworth model and the Gaussian copula model, which helps them better capture the extremal dependence structure; see equation (18) in Huser & Wadsworth (2019) for the censoring scheme. Then we fit the GH copula to the data without censoring low observations. Figure 5 displays the fitted values of 20-variate χ_u and η_u , as defined in Section 2.2, for the three models. The uncertainty measures are based on the same stationary




Figure 5: Estimates of 20-variate χ_u (left panel) and η_u (right panel) for the hindcast wave height data. Central black dots: empirical estimate of χ_u and η_u ; dashed line: approximate 95% confidence intervals based on a stationary bootstrap procedure; thick solid red line: fit from the GH copula; thin solid blue line: fit from the Huser-Wadsworth model; thin solid green line: fit from the Gaussian model.

bootstrap procedure as in Huser & Wadsworth (2019), i.e., 200 bootstrap samples are generated using the stationary bootstrap which sample blocks of geometric length. Strikingly, the results show that the GH copula can capture the extremal dependence better than the Gaussian copula, even though a censored likelihood has been used for the Gaussian copula. The performance of the GH copula is slightly worse than the Huser–Wadsworth copula, but the differences are fairly minor overall and the GH model still provides a very good fit in the upper tail despite its full uncensored likelihood estimation approach. We stress again that the Huser–Wadsworth copula is specially designed to model the upper tail and a censored likelihood has been used for this purpose, while our proposed GH model is here used to model the entire distribution, from low to high quantiles. We also fit the Huser–Wadsworth model and the Gaussian model without censoring the low observations. The log-likelihood of GH copula is around 2000 units larger than that of the Huser–Wadsworth model and 6000 units larger than that of the Gaussian copula model, which indicates that the GH copula is clearly more flexible than the other two if all observations are considered.

4.2 Wind Gust Example

To illustrate the computational benefits of our proposed model, we now consider a higher-dimensional dataset of daily maximum wind gusts from the state of Oklahoma, USA, which can be freely down-

Table 2: Number of parameters and AIC values for the different models fitted in our wind gust data application

Model	# of parameters	AIC
Gaussian copula	2	-298898
t copula	3	-324069
Hyperbolic copula	5	-316636
NIG copula	5	-327035
GH copula	6	-328141

loaded from mesonet.org. This dataset contains daily measurement of highest 3-second wind speed from January 1, 1995 to December 31, 2020, at 120 observational stations across the state of Oklahoma. To induce approximate stationarity, we only consider the summer observations, i.e., observations in July, August and September, which are often the highest in a year. As some observational stations have many missing values, we delete the ones with more than 100 missing values, resulting in 95 sites with 2392 observations at each site. The minimum distance between sites is around 12 km and the maximum distance is around 800 km. The time series at each site appear approximately stationary, hence we work directly with this dataset and transform all the marginals to standard uniform using the nonparametric approach based on ranks. Alternatively, we could use the semiparametric approach of Coles & Tawn (1991) as in the wave height example.

We then fit the GH copula model to the standardized dataset, as well as four of its subclasses and limiting models, namely the Gaussian copula, the t copula model, the hyperbolic copula and the NIG copula model, based on a full likelihood approach. The computational time for fitting the GH copula is less than three hours on a laptop, which is a significant speed-up compared to the censored Huser–Wadsworth model in such high dimensions. It takes less than ten minutes to fit the Gaussian copula and t copula due to the efficient computation of their quantile functions in R. Table 2 shows the number of parameters in each model and the Akaike information criterion (AIC) of the fitted models. The results indicate that the general GH copula has the best performance as expected, and the NIG submodel has the second best performance. The large difference between the AIC values of the GH copula model and its limiting cases, the widely used Gaussian copula model and the t copula model, implies that there is a clear advantage of using the GH copula model for this dataset despite its higher model complexity.

To assess how well our model captures the dependence both in the bulk and the tail of the data,

we consider Spearman’s rank correlation S_ρ and the bivariate extremal dependence measures χ_u and η_u described in Section 2.2. We first calculate the empirical estimates of S_ρ , χ_u and η_u with $u = 0.95$ from the data, and their model-based counterparts from the estimated model. Then we compute the bivariate kernel density estimators for pairs of observed empirical S_ρ and fitted S_ρ , pairs of observed empirical χ_u and fitted χ_u , and pairs of observed empirical η_u and fitted η_u , using the R package *ks*. Figure 6 depicts the contours at level 25%, 50% and 75% of these kernel density estimators. The results show that our model can capture strong dependencies (at short and moderate distances) both in the bulk and the tail very well, but it overestimates weak dependencies (at large distances). This is not very surprising as there are fewer pairs of sites that are weakly dependent (at long distances), so that strongly dependent pairs influence model estimation to a higher degree. In particular, the spatially-constant mixing variable R is constrained to provide a good fit at short distances, which simultaneously induces stronger dependencies across the entire spatial domain. This can also be seen from the discussion on the values of η for special cases of the GH distribution in Section 2.3. When the mixing variable R has very light tails ($\psi \rightarrow \infty$), as the distance between sites tends to infinity and their correlation $\rho \rightarrow 0$, the residual dependence coefficient η would converge to $1/\sqrt{2}$, which is much larger than $1/2$ and indicates that a strong dependence is still present in this case.

5 Discussion

The GH distribution has a long history, and found popularity in financial modeling. In this work we have re-investigated its tail dependence properties, pointed out the mistakes in its derivations in the literature, and gave a corrected description. Based on this result, we propose to use the GH copula for spatial extremes when asymptotic independence is present, which contributes to the spatial extremes literature, as existing models in this setting that are both flexible and can be fitted efficiently are scarce. We demonstrate the flexibility of this model both by a simulation study and two environmental data applications.

Unlike the asymptotic dependence case, where the scientific community broadly agrees that max-stable models should be fitted to block maxima, and generalized Pareto models should be fitted to threshold exceedances, there is no such consensus in the asymptotic independence case yet. Wadsworth & Tawn (2012) argue that it is often more natural, especially in the presence of asymptotic independence, to model the extremes of original events rather than site-wise maxima, and they suggest using a censored likelihood approach to better capture the extremal dependence. Here, we did not censor the low




Figure 6: Contours at level 25%, 50% and 75% of the bivariate kernel density estimators for pairs of observed empirical S_ρ and fitted S_ρ (left), pairs of observed empirical χ_u and fitted χ_u (middle), and pairs of observed empirical η_u and fitted η_u (right). A good fit should have density contours concentrated around the red dashed diagonal line.

observations when using the GH copula, because, as shown in the wave height data application, the non-censored GH copula can capture the tail dependence even better than the censored Gaussian copula and the high flexibility of the GH copula allows this full likelihood approach. Moreover, in this case we can gain significant computational efficiency and also capture dependence in the bulk, as shown in the wind gust data application.

One interesting research direction is to investigate the tail properties in the limiting case of the GH distribution when $\psi = 0$, which is an interesting distribution in its own as its marginal distribution has asymmetric tails, i.e., one has a power decay and the other has an exponential decay, and it reduces to the t distribution when the skewness parameter γ tends to zero. Unfortunately, the two different methods that we used in this work to derive the tail properties both fail in the limiting case when $\psi = 0$ and the skewness parameter is not zero. Hence, this remains an open problem. Another future research direction is to investigate the tail dependence properties of more general normal mean-variance mixtures when the mixing distribution is specified only through its tail behavior. As discussed in the wind gust data application, a weakness of the GH copula model is that independence cannot be captured as distance tends to infinity. This is a typical weakness for all random scale (and location) models. Hence, one can try to overcome this issue and improve the performance of the GH copula model

to better capture long-range weak dependence structures, for which the random partitioning approach of Morris et al. (2017) might be useful.

Acknowledgement

This publication is based upon work supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No. OSR-CRG2017-3434. We gratefully acknowledge Philip Jonathan of Shell Research for providing the wave height data analysed in Section 4.1.

References

- Abramowitz, M., & Stegun, I. A. (Eds.) (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. National Bureau of Standards, United States of America, 10 ed.
- Barndorff-Nielsen, O. E. (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society of London, Series A*, 353, 401–419.
- Barndorff-Nielsen, O. E. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics*, 24(1), 1–13.
- Blaesild, P., & Jensen, J. L. (1981). Multivariate distributions of hyperbolic type. *Statistical Distributions in Scientific Work*, 4, 45–66.
- Bortot, P., Coles, S., & Tawn, J. A. (2000). The multivariate gaussian tail model: an application to oceanographic data. *Journal of the Royal Statistical Society (Series C)*, 49(1), 31–49.
- Castruccio, S., Huser, R., & Genton, M. G. (2016). Higher-order composite likelihood inference for max-stable distributions and processes. *Journal of Computational and Graphical Statistics*, 25(4), 1212–1229.
- Coles, S., Heffernan, J., & Tawn, J. A. (1999). Dependence measures for extreme value analyses. *Extremes*, 2(4), 339–365.
- Coles, S., & Tawn, J. A. (1991). Modelling extreme multivariate events. *Journal of the Royal Statistical Society (Series B)*, 53, 377–392.

- Dagpunar, J. S. (1989). An easily implemented generalized inverse Gaussian generator. *Communications in Statistics - Simulation and Computation*, 18, 703–710.
- Davison, A. C., & Huser, R. (2015). Statistics of extremes. *The annual review of statistics and its applications*, 2, 203–235.
- Davison, A. C., Padoan, S. A., & Ribatet, M. (2012). Statistical modeling of spatial extremes. *Statistical Science*, 27(2), 161–186.
- de Haan, L. (1984). A spectral representation for max-stable processes. *The Annals of Probability*, 12(4), 1194–1204.
- de Haan, L., & Pereira, T. T. (2006). Spatial extremes: models for the stationary case. *The Annals of Statistics*, 34(1), 146–168.
- Dombry, C., Engelke, S., & Oesting, M. (2016). Exact simulation of max-stable process. *Biometrika*, 103(2), 303–317.
- Embrechts, P., McNeil, A., & Straumann, D. (2001). Correlation and dependence in risk management: properties and pitfalls. In M. Dempster, & H. Moffatt (Eds.) *Risk Management: Value at Risk and Beyond*, (pp. 176–223). Cambridge University Press.
- Engelke, S., Opitz, T., & Wadsworth, J. L. (2019). Extremal dependence of random scale constructions. *Extremes*, 22, 623–666.
- Ferreira, A., & de Haan, L. (2014). The generalized Pareto process; with a view towards application and simulation. *Bernoulli*, 20(4), 1717–1737.
- Huser, R., Opitz, T., & Thibaud, E. (2017). Bridging asymptotic independence and dependence in spatial extremes using gaussian scale mixtures. *Spatial Statistics*, 21, 166–186.
- Huser, R., & Wadsworth, J. L. (2019). Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 114(525), 434–444.
- Huser, R., & Wadsworth, J. L. (2020). Advances in statistical modeling of spatial extremes. *Wiley Interdisciplinary Reviews: Computational Statistics*, e1537.

- Jamalizadeh, A., & Balakrishnan, N. (2019). Conditional distributions of multivariate normal mean-variance mixtures. *Statistics and Probability Letters*, *145*, 312–316.
- Kabluchko, Z., Schlather, M., & de Haan, L. (2009). Stationary max-stable fields associated to negative definite functions. *The Annals of Probability*, *37*(5), 2042–2065.
- Krupskii, P., Huser, R., & Genton, M. G. (2018). Factor copula models for replicated spatial data. *Journal of the American Statistical Association*, *113*(521), 467–479.
- Le, P. D., Davison, A. C., Engelke, S., Leonard, M., & Westra, S. (2018). Dependence properties of spatial rainfall extremes and areal reduction factors. *Journal of Hydrology*, *565*, 711–719.
- Ledford, A. W., & Tawn, J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, *83*(1), 169–187.
- Manner, H., & Segers, J. (2011). Tails of correlation mixtures of elliptical copulas. *Insurance: Mathematics and Economics*, *48*, 153–160.
- McNeil, A. J., Frey, R., & Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.
- Morris, S. A., Reich, B. J., Thibaud, E., & Cooley, D. (2017). A space-time skew- t model for threshold exceedances. *Biometrics*, *73*, 749–758.
- Nolde, N. (2014). Geometric interpretation of the residual dependence coefficient. *Journal of Multivariate Analysis*, *123*, 85–95.
- Nolde, N., & Wadsworth, J. L. (2020). Connections between representations for multivariate extremes. Available from <https://arxiv.org/abs/2012.00990>.
- Padoan, S. A., Ribatet, M., & Sisson, S. A. (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, *105*(489), 263–277.
- Prause, K. (1999). *The Generalized Hyperbolic Model: Estimation, Financial Derivatives, and Risk Measures*. Ph.D. thesis, Albert-Ludwigs-Universität Freiburg.
- R Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.

- Rootzén, H., & Tajvidi, N. (2006). Multivariate generalized Pareto distributions. *Bernoulli*, *12*(5), 917–930.
- Schlueter, S., & Fischer, M. (2012). The weak tail dependence coefficient of the elliptical generalized hyperbolic distribution. *Extremes*, *15*, 159–174.
- Schmidt, R. (2003). *Dependencies of Extreme Events in Finance*. Ph.D. thesis, University of Ulm.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'institut de Statistique de l'Université de Paris*, *8*, 229–231.
- von Hammerstein, E. A. (2016). Tail behavior and tail dependence of generalized hyperbolic distributions. In J. Kallsen, & A. Papapantoleon (Eds.) *Advanced Modelling in Mathematical Finance*, (pp. 3–40). Cham: Springer.
- Wadsworth, J. L., & Tawn, J. A. (2012). Dependence modelling for spatial extremes. *Biometrika*, *99*(2), 253–272.
- Wadsworth, J. L., & Tawn, J. A. (2019). Higher-dimensional spatial extremes via single-site conditioning. Available from <https://arxiv.org/abs/1912.06560>.